The Linguistic Sophistication of Morphological Decomposition

The Linguistic Sophistication of Morphological Decomposition:

More than Islands of Regularity

By

Roberto Petrosino

Cambridge Scholars Publishing

The Linguistic Sophistication of Morphological Decomposition:
More than Islands of Regularity

By Roberto Petrosino

This book first published 2024

Cambridge Scholars Publishing

Lady Stephenson Library, Newcastle upon Tyne, NE6 2PA, UK

British Library Cataloguing in Publication Data
A catalogue record for this book is available from the British Library

ISBN (10): 1-5275-8614-6
ISBN (13): 978-1-5275-8614-7

Contents

List of Figures

List of Tables

catholic, *n.*
A person addicted to cats.

retired, *adj.*
Tired again.

fasten, *v.*
Make or become fast.

— *Decomposition matters*. And you may never realize it.

Chapter 1

Morphological decomposition: what, when, how

1.1 Decomposition and the lexicon

In the past few decades, a wealth of evidence has suggested that before being recognized, words are decomposed into smaller units that seem to correspond to morphemes (among others, Rastle et al., 2000, 2004; Fiorentino and Poeppel, 2007; Solomyak and Marantz, 2010; Lewis et al., 2011). This procedure, called MORPHOLOGICAL DECOMPOSITION, has been found to adhere to the following three-way pattern: (i) it occurs in words that are made of more than one morpheme (morphologically transparent words; e.g., *driver→{drive}-{er}*, meaning "someone who drives"); (ii) it occurs in words that only appear to be made of more than one morpheme, but are actually monomorphemic (morphologically opaque words; e.g., *brother→{broth}-{er}*, even if it does not mean "someone who broths"); (iii) it does not occur in words that contain a root plus additional phonological material devoid of any morphological information (e.g., *brothel↛{broth}-{el}*, where *el* is not an English suffix). This pattern of results suggests that the morpho-orthographic form of morphemes (that is, the phonological realization and the orthographic sequence of letter strings associated with a given morpheme) drives decomposition (since *-er* elicits decomposition, but *-el* does not), whereas semantics does not seem to matter (since *-er* elicits decomposition even in morphologically opaque, monomorphemic words like *brother*). This book builds on these findings and asks, in addition to the morpho-orthographic form of morphemes, what other linguistic properties affect decomposition (Figure 1.1, dotted section of the rectangle).

(1) EMPIRICAL QUESTION OF DECOMPOSITION
In addition to the morpho-orthographic forms of the morphemes, what linguistic properties, if any, affect decomposition?

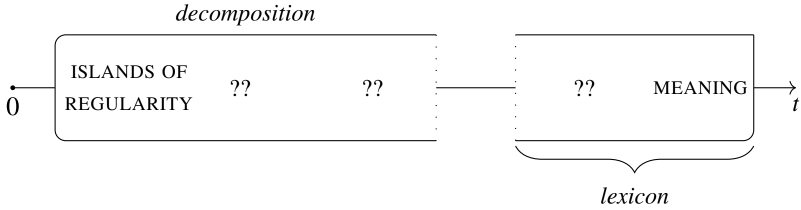

Figure 1.1: Time-course of the unfolding of word information in word processing (initial version). After being presented (0), a word is processed though a series of stages unfolding over time *t*. At each stage, a given set of properties is assumed to be available for processing purposes. *Decomposition* reportedly accesses the form of the constituent morphemes of a word via an orthography-based segmentation procedure, before contacting the *lexicon* (i.e., before accessing, for example, meaning). The '??' markings are provisional placeholders for the properties that will be tested in the upcoming chapters.

In asking (1), this book indirectly provides insights regarding the long-standing debate on the organization of lexical properties during word recognition. In the traditional view, all lexical properties are contained in a single place–the *mental lexicon* (2). Under the modular view, lexical access is all-or-none: if one property is accessed, all of other properties are accessed, too. However, this prediction is challenged by the psycholinguistic evidence showing that decomposition is sensitive to the orthographic form of morphemes, but not to their semantic properties (see sec. 1.3 for further details).

(2) MODULAR MENTAL LEXICON
All lexical and linguistic properties associated to morphemes are stored in one place–the *mental lexicon*.

To accommodate (2), psycholinguistic models of lexical access have argued that decomposition is not sensitive to the actual stored morpho-orthographic form of morphemes; rather, it is sensitive to statistical regularities in letter strings that are part of the reader's knowledge ("islands of regularity": Rastle and Davis, 2008; see sec. 1.4.2). For example, the string $ e r $ occurs very often in the English lexicon, and is therefore much likelier to form a morpheme than the string \$ e l \$.[1] In this sense, the former triggers decomposition, while the latter does not. This perspective is still compatible with the modular view of the mental lexicon (2) and, crucially, leads to the hypothesis (3), whereby decomposition occurs <u>before</u> any undeniably lexical property could be accessed ("contact with the lexicon"). As of today, (3) is commonly accepted, and also corroborated by recent neuromagnetic evidence (among others, Fruchter and Marantz, 2015).

(3) MORPHO-ORTHOGRAPHIC DECOMPOSITION HYPOTHESIS
Morphological decomposition relies on morpho-orthographic regularities and occurs <u>before</u> any contact with lexicon.

Admittedly, under (3), the question in (1) is theoretically vacuous; if decomposition occurs before any contact with the lexicon, the set of linguistic properties potentially affecting it must then be rather small, if not empty. However, as we will see in sec. 1.4.2, current models of decomposition do not strictly comply with (2). These models, more or less explicitly, seem to actually adopt a *distributed*, rather than modular, view of the mental lexicon, whereby availability of linguistic properties of morphemes is spread out across different stages of the recognition process. Therefore, the distributed view of the mental lexicon, as formalized in (4), validates the investigation reported here, while still assuming that decomposition relies on the aforementioned orthographic regularities.

(4) DISTRIBUTED MENTAL LEXICON
Lexical and linguistic properties associated with morphemes are accessed at different stages of the word recognition process.

In this sense, while primarily addressing the question in (1), this book aims to (i) situate each property tested in models of lexical access, and, indirectly, (ii) situate decomposition within word processing. To this end, I test the sensitivity of decomposition to a representative subset of higher-level linguistic properties. I primarily use the masked priming design, the primary experimental tool to probe morphological decomposition. In addition, I focus on decomposition in the visual modality, under the assumption that decompo-

[1] To ensure clarity, in this book I will use square brackets (e.g., [...]) to refer to the phonological forms of morphemes, and dollar signs (e.g., \$ . . . \$) to refer to the orthographic forms of morphemes.

sition may recruit different mechanisms in the auditory modality. The results I report here seem to be suggestive of a decomposition procedure that is more complex than previously thought. First, in contrast with what is currently assumed, decomposition does not seem to (at least solely) rely on morpho-orthographic regularities (see Chapter 4). Second, the recognition of decomposed morphemes seems to be sensitive to the lexicality status of the whole-word visual string (i.e., whether or not it forms a real word; see Chapter 2) and whole-word frequency (i.e., *dominance*; see Chapter 3), but not to syntactic affixal restrictions (see Chapter 5). I take stock of the whole investigation in Chapter 6, where I discuss the theoretical implications of the results obtained. In particular, these results seem to challenge the modular view (2) in favor a distributed view (4), in which some properties may be accessed at early stages of processing (thus affecting decomposition), and some other properties may only be available at later stages. I claim that such a seemingly clear divide is due to the time needed for the different properties to be properly accessed and computed. A competition-based model of decomposition is therefore proposed, in which multiple decomposition pattern candidates of the same visual stimulus are evaluated in parallel in order to identify the most likely pattern to decompose the stimulus.

This chapter is meant to give a concise, but hopefully thorough introduction to the literature on visual morphological decomposition and, more generally, visual word processing. Section 1.2 reviews the features of the priming design. Section 1.3 summarizes the main results from the priming literature on visual morphological decomposition. Section 1.4.2 reviews four major models of lexical access; these models will be used to interpret the results of the experiments reported in the following chapters. Section 1.5 describes the general experimental methodology adopted in the book. Finally, section 1.6 briefly outlines the following chapters.

1.2 The priming design

This book primarily uses the lexical decision priming paradigm. Generally speaking, the priming paradigm consists of two stimuli being presented one after the other; the first stimulus is called *prime*, the second is called *target*. The experimental task always concerns the target stimulus; in this book, the main task will always be a lexical decision task, in which subjects are asked to decide whether the target stimulus is an existing word by pressing a response button box. It has been shown that recognition of the target stimulus is faster when it is related to the prime stimulus to some degree: e.g., the target word *cat* is recognized as a word faster when it is preceded by a semantically related prime word like *kitten* than when it is preceded by an unrelated prime word like *hand*.

Priming experiments typically have two main conditions. A RELATED CONDITION, in which the prime and the target are related along some dimension (e.g., they may be identical: *cat-CAT*; orthographically similar: *cap-CAT*; or morphologically related: *cats-CAT*) and an UNRELATED CONDITION (baseline condition), in which the prime and the target are completely different and, crucially, unrelated, items (e.g., *hand-CAT*).[1] There are two main designs in which the related and unrelated conditions may be arranged. The basic design is the *Latin Square* design. Latin Square designs involve splitting the whole set of target stimuli that are presented (e.g., words) in the number of conditions to be tested (e.g., related and unrelated conditions), and rotating the groups of targets through the conditions to create different *lists*. Subjects are presented with one list only, so that they see exactly the same targets (which may be preceded by a related or unrelated prime), regardless of the list they were assigned. This kind of design has the advantage of naturally controlling for stimulus properties (e.g., frequency, orthographic/morphological/phonological length, neighborhood density, etc.).[2] On the other hand, Latin Square designs are not often used in lexical decision tasks because they are only possible when all target stimuli can be paired with all prime stimuli across conditions. This is rare, and mostly happens when the primes are pseudo-words. The standard design that is typically used instead involves creating each related condition separately (e.g., morphological condition: *driver-DRIVE*) and then arranging the stimuli across two different lists, such that, in each list, half of the targets of each condition are preceded by the corresponding related primes and the other half are preceded by the corresponding unrelated primes. This way, each list consists of several self-contained sub-designs; in each sub-design, the targets are either related to the prime (related condition: *driver-DRIVE*) or unrelated to the prime (unrelated condition: *lovely-DRIVE*). While this kind of design offers more freedom (because it allows for counterbalancing across word-word

[1] In visual priming experiments, targets are commonly presented in UPPERCASE and primes in lowercase. This is to ensure that lexical decisions are not influenced by merely graphical similarity between the prime and the target. For the sake of clarity, I will maintain this typographical difference between primes and target throughout the entire book.

[2] It should be noted that the actual extent to which lexical information of words may actually affect priming effects has not yet been shown. However, this is common practice in the field, as a way to remove any potential confounding effect resulting from it.

pairs, unlike Latin Square designs), it has the drawback of not naturally controlling for the lexical properties of the stimuli. This means that primes and targets of all conditions must share a similar distribution of any given lexical information (e.g., frequency or length) across all conditions, in order to guarantee similar distributions (but see fn. 2).

Once response times (RTs) to target recognition have been recorded, they are averaged together and used to calculate priming effects for each condition. *Priming effects* are calculated as the difference between the averaged RT to the related condition and the averaged RT to the unrelated (i.e., baseline) condition. According to the most common practice in priming studies, the sign of the difference is then reversed for clarity.

(5) $$priming = -(RT_{related} - RT_{unrelated})$$

If positive, the priming magnitude calculated in (5) expresses the amount of *facilitation* (expressed in ms) that the presentation of a related prime exerted on the recognition of the corresponding target. If negative, the priming magnitude calculated in (5) expresses, instead, the amount of *inhibition* (expressed in ms) that the presentation of a related prime exerted on the recognition of the corresponding target.

Morphological decomposition has mainly been investigated by adopting a sub-type of the visual priming design. This design is called *masked* priming and was pioneered by Forster and Davis (1984). In this design, a *mask* (that is, a series of hashmarks, #####) is presented for 500 ms before the prime stimulus, which is usually presented for 30-70 ms. At such a short duration, primes are not consciously recognizable to subjects. The advantage of the masked priming technique is two-fold: (i) it prevents facilitation effects due to strategic processes triggered by episodic memory and (ii) it elicits facilitation effects before full lexical access of the prime word is complete. The target immediately follows the prime and usually stays on the screen until the subject performs the task on the target (e.g., lexical decision). As already discussed in sec. 1.1, I will follow the literature on morphological decomposition and use the masked priming design as the primary tool to explore our research questions. In doing this, I consider the masked priming design as the right tool to probe the linguistic sophistication of early decomposition, under the assumption that prime masking prevents the prime stimulus from being further processed after target presentation. As we will see in Chapter 6, this methodological assumption about the experimental design used is crucial for the results to be properly interpreted.

1.3 Morphological decomposition as revealed in visual masked priming: a review

In this section, I briefly summarize the main results on morphological decomposition as reported in the visual priming literature. The review below represents the bulk of the findings I will be referring to throughout the book. Table 1.1 on page 9 is meant to serve as a summary of the relevant results the reader can go back to whenever needed.

1.3.1 Root priming

Priming experiments on morphological processing have primarily focused on root priming, in which the prime is the corresponding bare root of the correspondent target (e.g., *driver-DRIVE*). The studies by Rastle et al. (2000) and Rastle et al. (2004) were among the first to report root priming in a masked design. Rastle et al. (2000) looked at priming effects in four different conditions: in the morphological condition, the prime derives from the target (e.g., *adapter-ADAPT*); in the semantic condition, the prime and the target are semantically related (e.g., *cello-VIOLIN*); in the identity condition, the prime and the target are the same word (e.g., *church-CHURCH*); finally, in the orthographic condition, the prime shares the first letters with the target (e.g., *electrode-ELECT*). The four conditions were tested across three different prime-target SOA (Stimulus Onset Asynchrony, namely the time between the presentation of the prime and the target) durations. At SOAs of 42 and 72 ms, pairs in both the morphological and identity conditions showed priming effects, whereas pairs in the semantic and the orthographic condition did not. At SOA of 230 ms, pairs in the identity, morphological, and semantic conditions showed priming effects, while pairs in the orthographic condition did not. The fact that unlike semantic priming, morphological priming arises even at short SOAs has been interpreted as confirmation of the hypothesis that morphological decomposition occurs subliminally (namely, even at short SOAs, in which the prime duration is so fast that the presentation of the prime before the target cannot be acknowledged by subjects) and cannot be induced by semantic relatedness (since semantically related, but morphologically unrelated pairs such as *cello-VIOLIN* did not show priming at short SOAs). Building on these results, Rastle et al. (2004) looked at subliminal priming elicited in the following three conditions: the morphologically transparent condition, in which prime and target words were morphologically, semantically and orthographically related (*alarming-ALARM*); the morphologically opaque condition, in which prime and target words looked like the former was a derived form of the latter, but were not semantically related (*brother-BROTH*); finally, the orthographic condition, in which prime and target words were only orthographically related so that the target word was contained in the prime word (*brothel-BROTH*). With a 42ms-long SOA, priming effects were found in the transparent and opaque conditions, and not for the orthographic condition. Taken together with Rastle et al. (2000)'s results, these results further suggest that decomposition occurs only in the presence of real morphemes, since *brother* primes *BROTH* as much as *alarming* primes *ALARM* (even though *brother* and *broth* are not semantically related to one another at the synchronic level), but *brothel* does not prime *BROTH* (possibly because *el* is not a real English morpheme).

The studies that were subsequently published can all be seen as aiming to further understand the mechanism of decomposition. Morris et al. (2011) compared the priming effects in three conditions: a morphological condition (*flexible-FLEX*); a pseudo-morphological condition, in which the pseudo-word prime consisted of the corresponding monomorphemic target suffixed with a real, but syntactically illicit morpheme (e.g., *flexity-FLEX*); and a pseudo-affixed condition, in which the pseudo-word prime consisted of the corresponding monomorphemic target suffixed with a phonologically licit, non-suffixal ending (e.g., *flexire-FLEX*). All three conditions were found to elicit priming effects whose magnitude did not vary across conditions. These results are at odds with the results reported in Rastle et al. (2000, 2004); in particular, the pseudo-affixed condition eliciting priming is at issue with the argument that decomposition occurs only when the prime contains extant morphemes. I return to this apparent contradiction below (sec. 1.3.3).

1.3.2 Affix priming

Only a few studies have looked at affix priming. Dominguez et al. (2010) tested priming effects in response to Spanish prefixes by comparing two conditions: a prefixed condition, in which prime and target shared the same prefix (e.g., *infeliz-INCAPAZ* 'unhappy-incapable') and a syllabic condition, in which prime words had the first syllable being phonologically and orthographically identical to the prefix of the corresponding target (*industria-INCAPAZ* 'industry-incapable'). In line with Rastle et al. (2004), Dominguez et al. (2010) found significant priming effects for both the prefixed and the syllabic condition only at the short SOA of 33 ms (experiment

PRIMING TYPE	CONDITION	EXAMPLE	PRIME LEXICALITY	PRIMING?	REFERENCE
ROOT PRIMING	transparent	*driver-DRIVE*	W	Y	Rastle et al. 2000, 2004
	opaque	*brother-BROTH*	W	Y	Rastle et al. 2004
	root-containing	*brothel-BROTH*	W	N	
	root+suffix	*flexity-FLEX*	NW	Y	Morris et al. 2011
	root+non-suffixal ending	*flexire-FLEX*	NW	Y	
PREFIX PRIMING	pseudo-prefixed	*industria-INCAPAZ*	W	Y	Dominguez et al. 2010
	prefix+root	*infeliz-INCAPAZ*	W	Y	
	transparent	*impatient-IMMOBILE*	W	Y	Chateau et al. 2012
	opaque	*imitate-IMMOBILE*	W	Y	
	orthographic	*pursuit-PURPLE*	W	N	
SUFFIX PRIMING	suffix only	*dad-IGUALDAD*	NW	Y	Duñabeitia et al. 2008
	pseudo-suffix only	*men-CERTAMEN*	NW	N	
	transparent	*brevidad-IGUALDAD*	W	Y	
	orthographic	*volumen-CERTAMEN*	W	N	
	suffix	*sheeter-TEACHER*	NW	Y	Crepaldi et al. 2015
	word-ending	*poolel-BARREL*	NW	N	
COMPOUND PRIMING	non-head constituent	*flagpole-FLAG*	W	Y	Fiorentino & Fund-Reznicek 2009
		hallmark-HALL	W	Y	
		plankton-PLAN	W	N	
	head-final constituent	*classroom-CLASS*	W	Y	
		honeymoon-MOON	W	Y	
		battalion-LION	W	N	
	pseudo-word+word	*slegrack-RACK*	NW	Y	Fiorentino et al. 2015
	word+word	*drugrack-RACK*	NW	Y	

Table 1.1: Summary of results of visual priming studies. *Legend:* W, word; NW, non-word.

1a); the effects to the syllabic condition gradually decreased as the SOA increased, with the effects to the prefixed condition not varying throughout (experiment 1b, c). In English, Chateau et al. (2002) found priming effects for both transparently prefixed (*impatient-IMMOBILE*) and opaquely prefixed (*imitate-IMMOBILE*) pairs, but no priming effects for orthographic pairs sharing the initial letters (*pursuit-PURPLE*). Moreover, transparent pairs elicited priming regardless of whether the prefix was high or low in spelling-meaning consistency (for example, *im-* is high consistency, whereas *con-* is low consistency), whereas opaque pairs elicited priming only when affixed with high-consistency prefixes (e.g., *context-CONFORM* did not prime).

As for suffix priming, Andoni Duñabeitia et al. (2008) reported a series of Spanish experiments comparing a morphological condition, in which prime and target stimuli shared the same suffix and an orthographic condition in which prime and target stimuli shared the same last three letters. In the morphological condition, the target words (e.g., *IGUALDAD* 'equality') were preceded by primes being the sole corresponding suffixes (*dad*; experiment 1), the suffixes along with non-alphabetic symbols (*%%%%dad*; experiment 2), or words sharing the same suffix (e.g., *brevedad* brevity; experiment 3). Similarly, in the orthographic condition, the target words (e.g., *CERTAMEN* 'contest') were preceded by the primes that were only the corresponding pseudo-suffixes (*men*; experiment 1), the pseudo-suffixes attached to non-alphabetic symbols (i.e., *%%%%men*; experiment 2), or words sharing the same ending (e.g., *volumen*; experiment 3). In all three masked experiments, priming effects were found for all morphological conditions only, regardless of whether the suffixes were presented stand-alone (experiment 1), together with non-alphabetic symbols (experiment 2), or embedded in a real bimorphemic word (experiment 3). This suggests that decomposition occurs in the presence not only of alphabetic strings (e.g., *brevidad* is decomposed into *breve-dad*), but also of non-alphabetic strings (i.e., *%%%%dad* is decomposed into *%%%%-dad*). Finally, Crepaldi et al. (2016) is the only study I know of that looked at visual suffix priming in English. The RTs to pairs in which pseudo-word primes shared the same suffix with derived targets (suffix condition; *sheeter-TEACHER*) were found to be faster than the RTs to pairs in which pseudo-word primes were affixed with a different suffix (suffix control condition; *sheetal-TEACHER*) and the RTs to pairs in which pseudo-word primes were affixed with a non-suffixal ending (control condition; *sheetub-TEACHER*). On the other hand, the RT to pairs in which pseudo-word primes shared the same non-suffixal ending with monomorphemic targets (word-ending condition; *poolel-BARREL*) were found not to be different from the RTs to pairs in which pseudo-word primes had a different suffix (word-ending control condition; *poolic-BARREL*) or from the RT to pairs in which pseudo-word primes has a different non-suffixal ending (control condition; *poolut-BARREL*). This suggests that morphological priming may be elicited only when a given letter cluster constitutes a real morpheme.

1.3.3 Compound priming

A substantial body of studies suggests that decomposition occurs in compounds, too. In the priming literature, both compound-initial and compound-final constituents were reported to be primed by their respective compounds. Fiorentino and Fund-Reznicek (2009) looked at priming effects to both head and non-head compound constituents in three conditions: the transparent, opaque, and orthographic conditions. In experiment 1, targets were non-head constituents of both transparent and opaque compounds (e.g., transparent condition: *flagpole-FLAG*; opaque condition: *hallmark-HALL*). In experiment 2, targets were head constituents of the both transparent and opaque compounds: e.g., transparent condition: *classroom-ROOM*; opaque condition: *honeymoon-MOON*. In both experiments, the orthographic condition consisted of a pseudo-compound prime and the target contained therein (experiment 1: *plankton-PLAN*; experiment 2: *battalion-LION*). In both experiments, the transparent and opaque conditions elicited similar priming effects and the orthographic condition did not.

The apparent contradiction between the results reported in Rastle et al. (2004) and Morris et al. (2011; see sec.1.3.1) was first acknowledged in Fiorentino et al. (2015). Recall that in previous visual priming studies decomposition seems to occur in two apparently contradicting circumstances. It occurs in response to primes sharing a morphologically transparent (e.g., *alarming*) or opaque (e.g., *brother*) relationship with their targets Rastle et al. (e.g., *ALARM, BROTH*; 2004). It also occurs in response to primes that contain the corresponding target (*FLEX*), regardless of whether the primes are licit affixed words (*flexible*), illicit affixed pseudo-words (*flexity*), or pseudo-words ending with a non-suffixal letter cluster (*flexire:* Morris et al., 2011). Fiorentino et al. (2015) used two kinds of novel compounds to prime the corresponding compound-final constituent: compounds made of two extant English words (e.g, *drugrack-RACK*) and compounds made of an English pseudo-word and an extant English word (e.g., *slegrack-RACK*). The results showed priming effects in both conditions, in line with Morris et al. (2011), and challenging Rastle et al. (2004)'s argument whereby decomposition occurs only for extant morphemes. The next chapter will address this contradiction and try to find a way of resolving the contradiction while maintaining the hypothesis in (3).

1.4 From letters to words: Models of visual word recognition

While exploring the linguistic sophistication of decomposition, the book indirectly engages with (i) orthographic processing, i.e. the process whereby orthographic stimuli (i.e., letters) are decoded by the parser; and (ii) lexical access, i.e., the process whereby the linguistic and lexical properties of words are retrieved. This section is divided in two subsections. In the first, I describe the time-course of the orthographic word stimulus and briefly review the currently well-accepted model of orthographic processing; no other model will be taken into consideration, since the debate on orthographic processing is irrelevant to the purposes of this book. In the second subsection, I take a deeper look at how decomposition and lexical access are believed to occur in four different models; as being specifically relevant to our purposes, these models will be worked through in each of the following chapters.

1.4.1 Orthographic processing

When a visual word is presented, it is recognized (i.e., it is *read*) within a fraction of a second, regardless of the potential variations in position, size, case, and font (henceforth, *orthographic-feature invariance*). Invented less than 6,000 years ago, it is unlikely that reading (also known as visual word recognition) taps into a specific neuro-cognitive mechanism implemented in the brain. Rather, it is far more likely that reading recruits (or "recycles"; Dehaene and Cohen, 2007) general visual object identification neural mechanisms, which eventually trigger the development of letter-specific detectors due to the extensive training readers are exposed to from school age (Grainger, 2018). In particular, it has been shown that a localized region of the left occipito-temporal sulcus, close to the fusiform gyrus – commonly labeled as the Visual Word Form Area (VWFA) – is particularly sensitive to letter identification (Cohen et al., 2000).

In this book, I adopt the anatomically-implemented neurobiological framework proposed by Dehaene et al. (2005), which takes advantage of the large amount of evidence about the organization of the ventral occipito-temporal route for visual object recognition in the human brain. While other models argue that letter feature-invariance is achieved in a single step, Dehaene et al. (2005)'s hierarchical model of orthographic recognition argues instead that visual receptive fields gradually widen their visual scope, aligning with the increase in the degree of complexity and abstractness of the properties neuronal populations that are sensitive to. At lower levels, only basic features are detected (e.g., local contrasts). At the intermediate level, letter contours and (case-specific) shapes are used to detect a letter. At the higher level of analysis, the letter just detected is matched with the corresponding abstract (i.e., feature-invariant) letter representation. Once letters are abstractly identified, the system starts grouping letters together to form bigram sequences. Bigram sequences are argued to be the best trade-off letter sequence that allows the system to be flexible about letter location and order (both within the sequence and with respect to the flanking letters and sequences), as compared to smaller (monogram) or bigger (e.g., trigram) sequences. Finally, bigrams are analyzed together in combination, so that they may be identified as morpho-orthographic units (namely, orthographic units that correspond to the abstract morpheme units).

1.4.2 Models of lexical access and decomposition

As already discussed in sec. 1.1, the models of lexical access that have been proposed vary in the ontology of the mental lexicon. In some models, the mental lexicon is purely modular, in the sense that all linguistic properties (e.g., the morpho-orthographic form, alternations, morpho-syntactic properties, meaning) are stored in a single place (e.g., Hauk et al., 2006; Pulvermüller et al., 2001; Sereno et al., 1998). In other models, the mental lexicon is distributed throughout word processing, so that properties may gradually unfold over time, as word processing occurs (e.g., Dehaene et al., 2005; Vinckier et al., 2007). In this book, I consider four models: the bin model (also known as the entry-opening model; Forster and Davis, 1984; Forster, 1999), the race model (Schreuder and Baayen, 1995), the full decomposition model (Taft, 1994, 2004), and the morpho-orthographic model (Crepaldi et al., 2010). On the one hand, all four models seem to share a fairly distributed view of the mental lexicon. It is also important to keep in mind that all of these models assume that *segmentation* (i.e., the acquisition process whereby the parser learns morpho-orthographic units) is distinct from *decomposition* (i.e., the on-line process of unit identification). On the other hand, they differ from one another in the mechanisms and, therefore, in the time-course of lexical retrieval. In the subsections below, I give a detailed description of the four models. The structure of each subsection is the same. In the first paragraph, I outline the mechanisms. In the second paragraph, I examine how lexical properties are distributed across stages. Finally, in the third paragraph, I exemplify how the model captures Rastle et al. (2004)'s critical results (see sec. 1.3.1). Dealing with *visual* word recognition, all of these models will be referred to when discussing the results of the visual experiments reported in Chapters 2-5 below. In doing this, I indirectly assume that different modality-specific strategies may be recruited in word processing (for a review, Grainger, 2018).[3]

A. BIN MODEL

Forster (1999) argues that orthographically similar words are grouped together in *bins*; each bin is identified with a specific hash code. When the stimulus is presented, the parser first converts the orthographic strings into a hash code indicating the bin where the correct entry is likely to be located. Recognition is then carried out through two stages. In the first stage – the entry-opening stage (originally proposed by Forster and Davis, 1984) – each entry within the same bin goes through a fast, frequency-ordered comparison with the stimulus and is assigned a goodness-of-fit score. During this stage, entries are sorted into subcategories depending on the degree of the matching comparison: perfect matches involve entries that are no different from the stimulus, close matches involve entries that are a bit different from the stimulus, and no-matches involve entries that are completely different from the the stimulus. Entries that are either perfect or close matches with the stimulus are flagged as potential candidates, and are therefore "opened," so that the information therein may be retrieved. The whole process operates in parallel for all entries and all identified bins. As soon as an entry is opened, the second stage–the verification stage–begins; at this stage, the comparison between the stimulus and the opened entries is much slower and detailed, and goes through each candidate in sequence. In this stage, if a candidate entry does not match with the stimulus, it is rejected, so that the verification of the next candidate begins. When a candidate entry is found to match the stimulus, the entry is selected as the correct one and all other entries are immediately closed even if they have not yet been evaluated. Priming is argued to arise when the entry for the

[3]The literature on modality-specific lexical access is vast and touches on several related topics, among which are phonological processing (Vitevitch et al., 1999; Vitevitch and Luce, 1998, 1999), syntactic processing (Dikker et al., 2009), and connectionist modeling (e.g., interactive activation models: McClelland and Rumelhart, 1981; parallel distributed models: Seidenberg and McClelland, 1989; the cohort model: Marslen-Wilson and Welsh, 1978; Taft and Hambly, 1986; the TRACE model: McClelland and Elman, 1986).